
The
Network Weather Service
and the
Information Power Grid

Rich Wolski

University of Tennessee

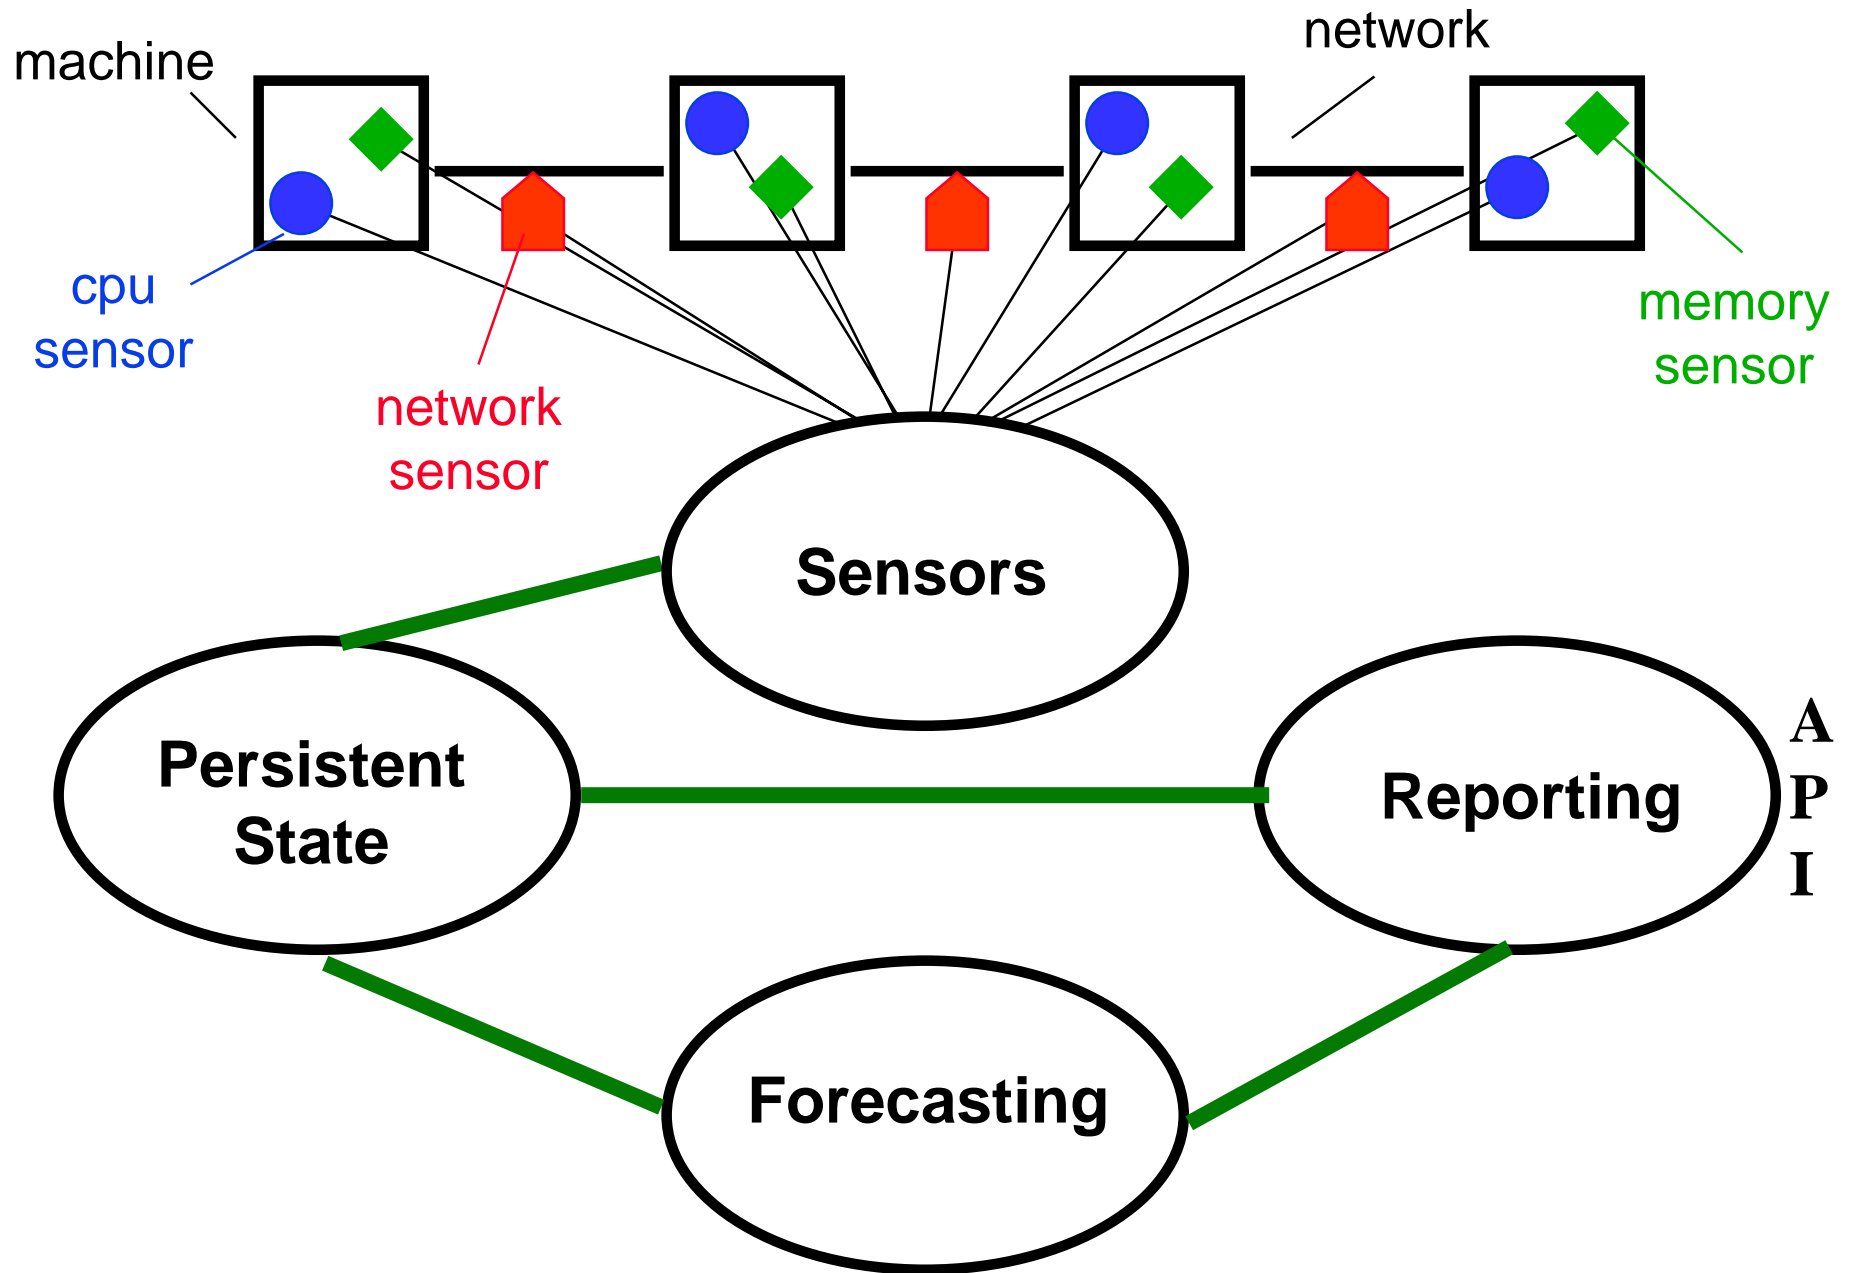
Resource Performance Monitoring and Prediction

- Heterogeneity: Deliverable performance depends on
 - Type of **hardware**
 - Installed **software** base
 - Local **management policies**
- Dynamism: Performance **varies over time**
 - Contention
 - Resource failure
 - Reconfiguration
- **Time-sensitive performance prediction**
 - *Scheduling*: current load is only interesting if it predicts what the load will be when the resource is used
 - *Fault diagnosis*: what is a fault and what is noise?

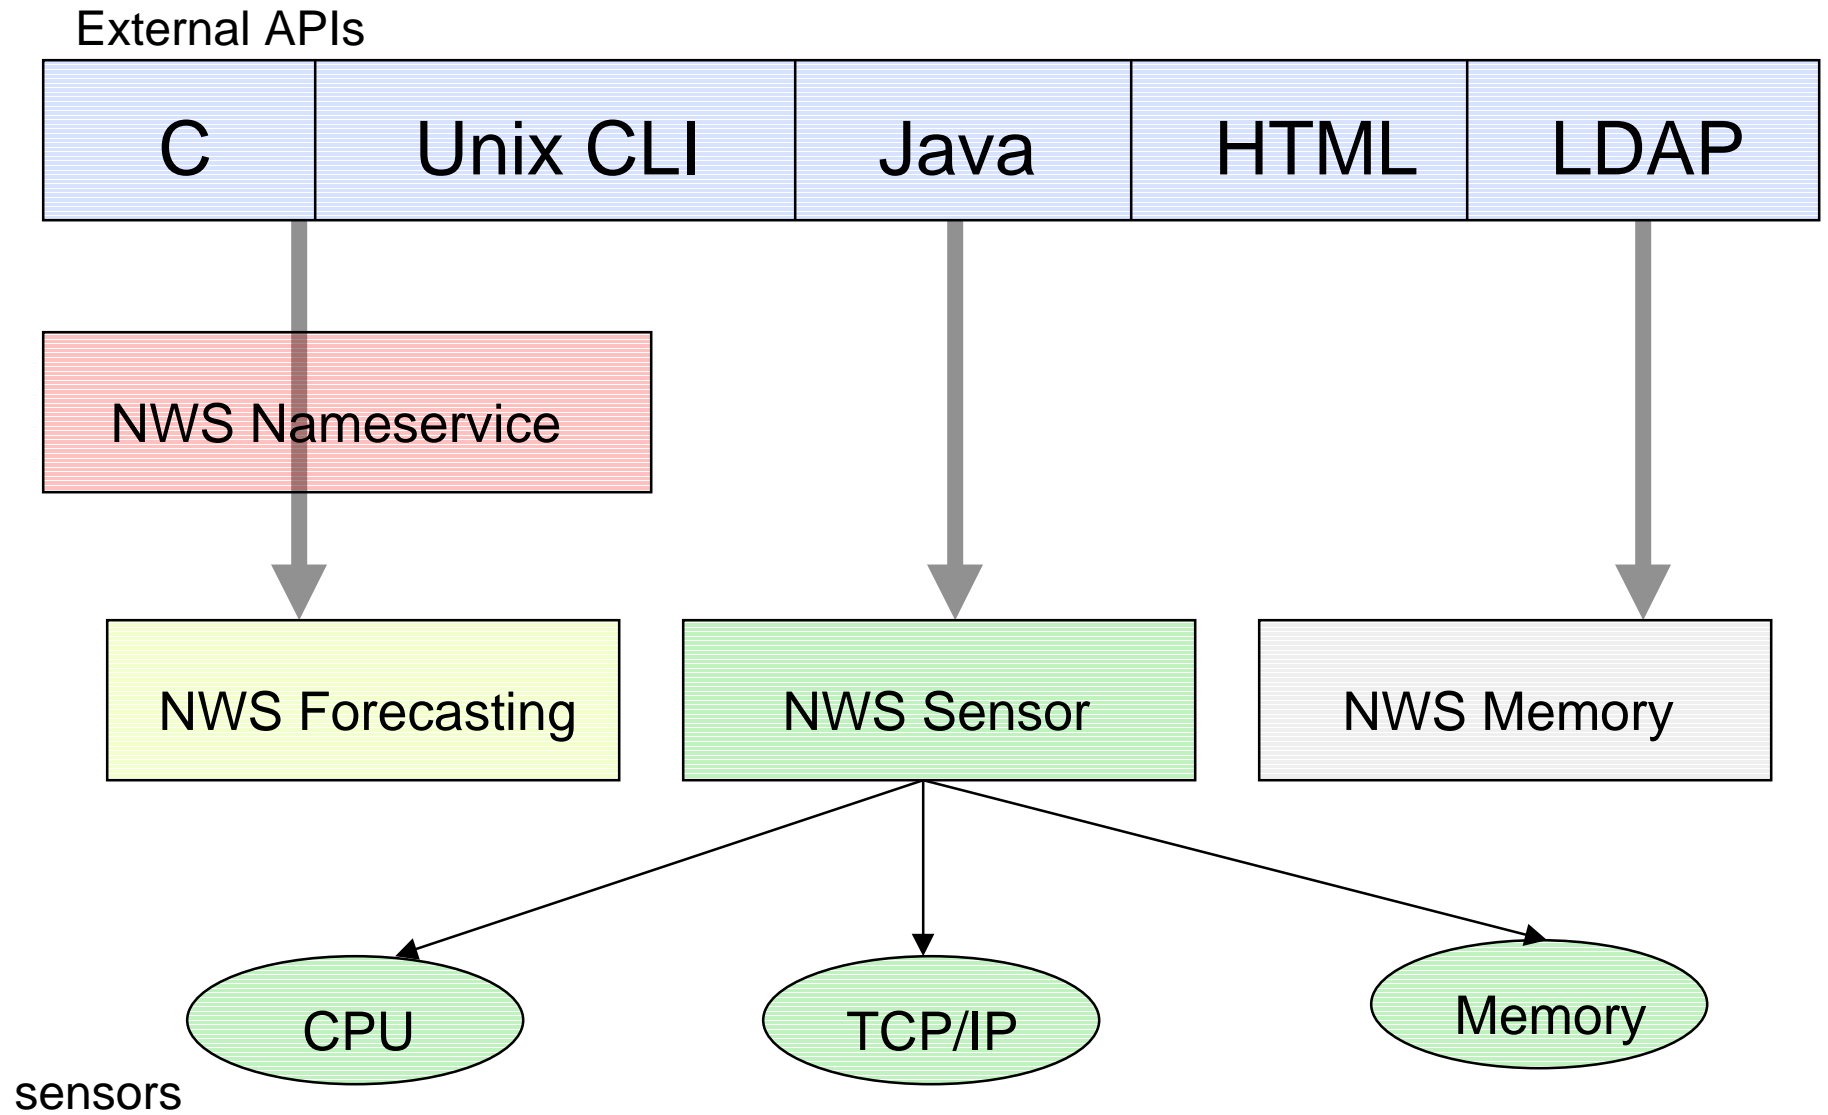
Network Weather Service Functionality

- The Network Weather Service (**NWS**)
 - **Monitors** the deliverable performance available from a distributed resource set.
 - **Forecasts** future performance levels using statistical forecasting models.
 - **Reports** monitor and forecast data to interested client schedulers, applications, visual interfaces, etc.
- **Generally available Grid service**
 - portable, extensible, robust, scalable, etc-able

Logical Architecture



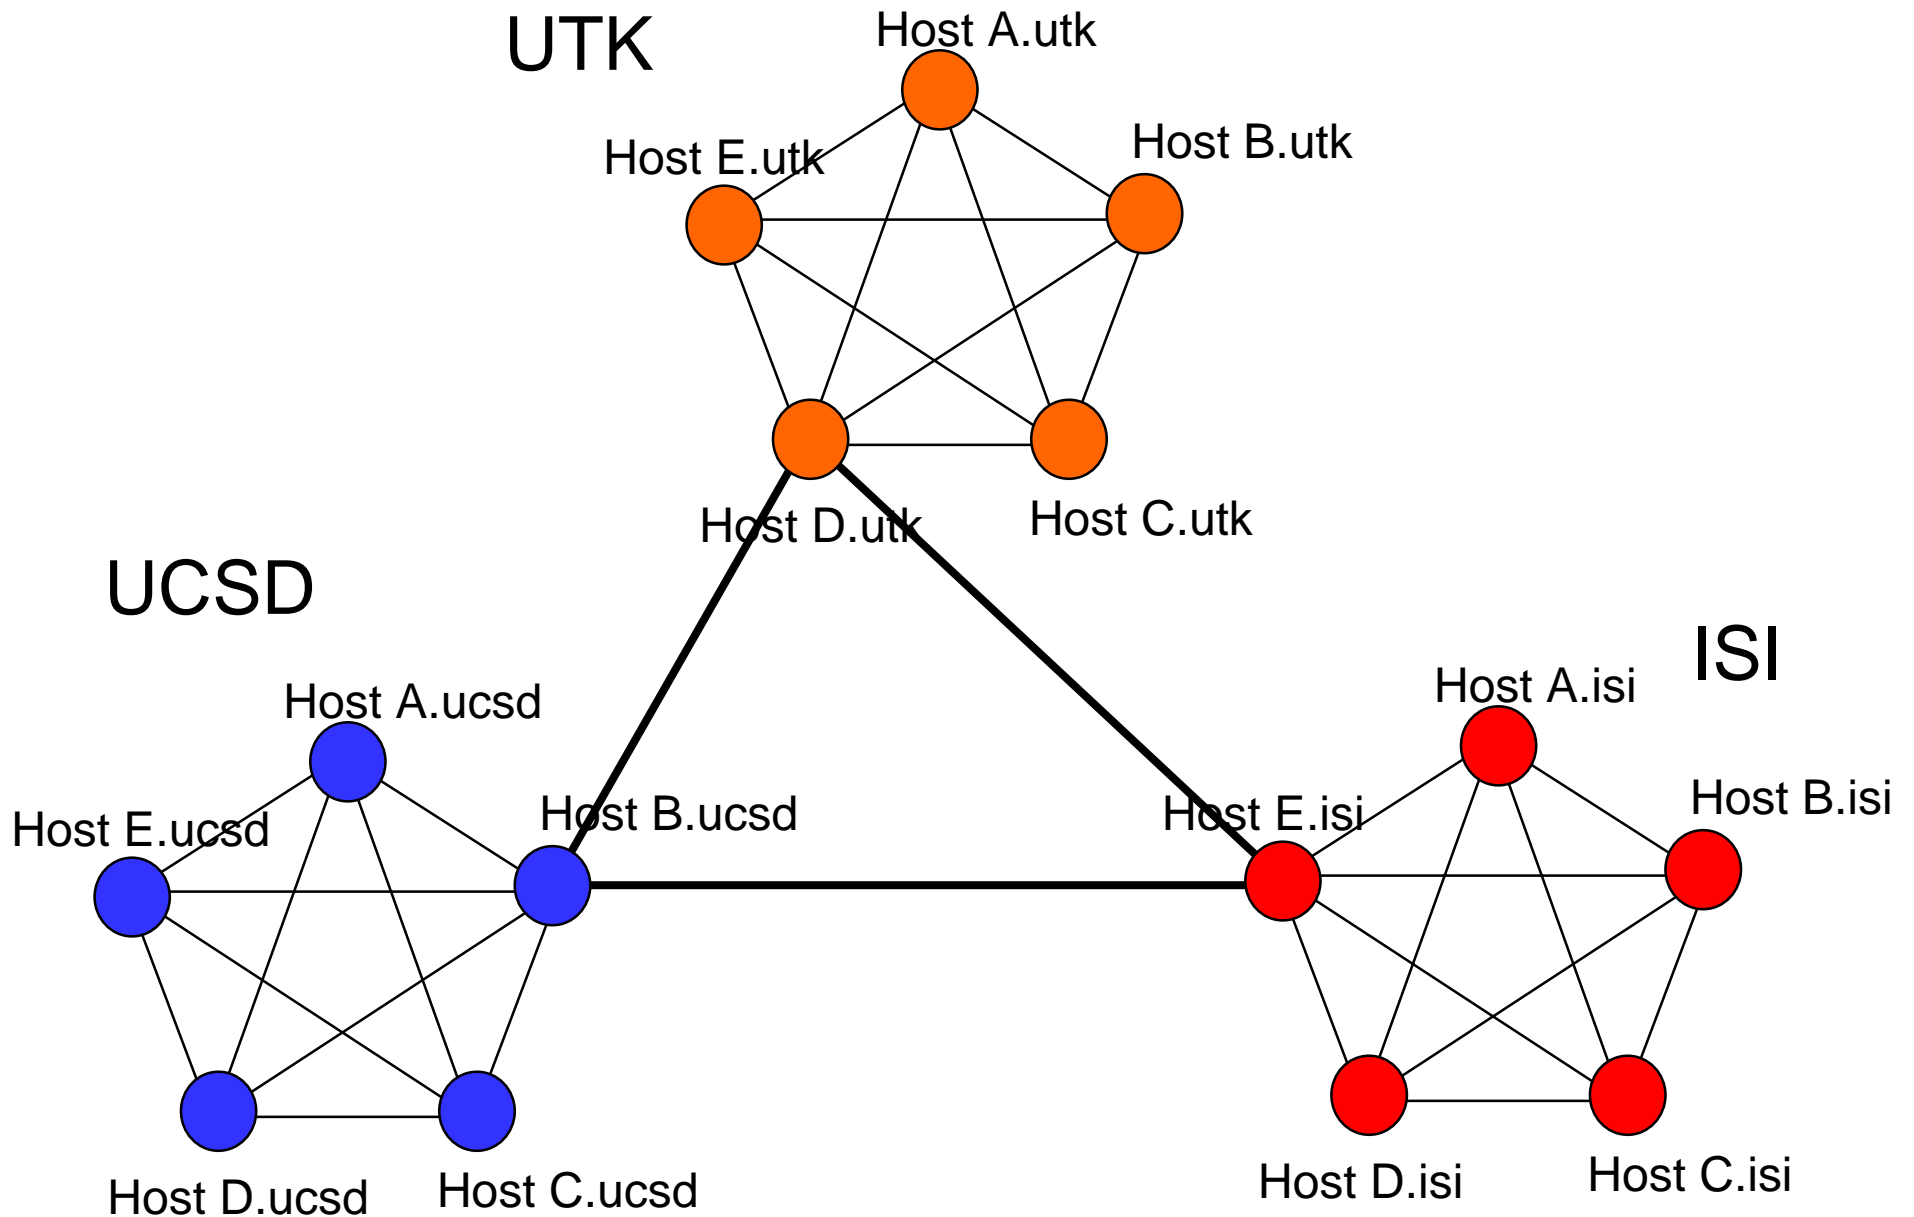
Software Architecture



Currently Supported Performance Sensors

- CPU sensor
 - Adaptively combines active process with load average
- Real memory sensor
 - Examines RSS of all executing processes to determine available real memory
- TCP/IP bandwidth and latency
 - Configurable buffer and message sizes
 - End-to-end measurements
- New sensors can be added via a well-defined interface
 - Periodic and synchronized sensor control

Scalability: Hierarchical Organization




Forecasting

- Each sensor produces a **time series** of measurements
- Forecasting modules apply **fast statistical** models to make short-term forecasts
 - Mean based forecasts (likelihood)
 - Median based forecasts (noise filtering)
 - Autoregression
- The system **dynamically chooses** the best model based on cumulative error
- New forecasters can be added via a **well-defined API**

NASA IPG and the NWS

- NWS Interface to Globus
- NWS application forecasting library – *dynamic benchmarking*
- Archival storage facility based on Netlogger
- AppLeS enhancements, support, and good will


The NWS and The Grid Information Service



- **NWS** provides network performance information on registered hosts on-demand
- Shell backend converts ldap requests to **NWS command-line invocations**
- Registration and clique configuration performed using **NWS** tools

Intended Network Configuration

Daemon and Clique Layout



Accessing the NWS via Globus

- To obtain network performance information, use the command
 - `grid-info-search -b "service=NWS, o=Grid" "<search-filter>"`
 - `grid-info-search -h nws.cs.utk.edu -p 390 /
-b "service=NWS, o=Grid" "<search-filter>"`
- Sample search filters
 - `source-hostname=bolas.isi.edu`
 - `destination-hostname=nws.cs.utk.edu`
 - `(&(source-hostname=pitcairn.mcs.anl.gov)
(destination-hostname=burns.ecs.csun.edu))`

Dynamic Benchmarking

- **NWS** Forecasting library can be loaded with the application
 - Maintains time series information internally
 - Memory footprint is tunable
 - Forecasting suite is configurable
- “Standard” NWS Forecasting suite requires 240 usec/forecast on a 233MHZ Pentium.
- API:

```
handle = INITFORECASTER( );  
UPDATEFORECASTER(handle,timestamp,measurement);  
forecast = FORECAST(handle);
```

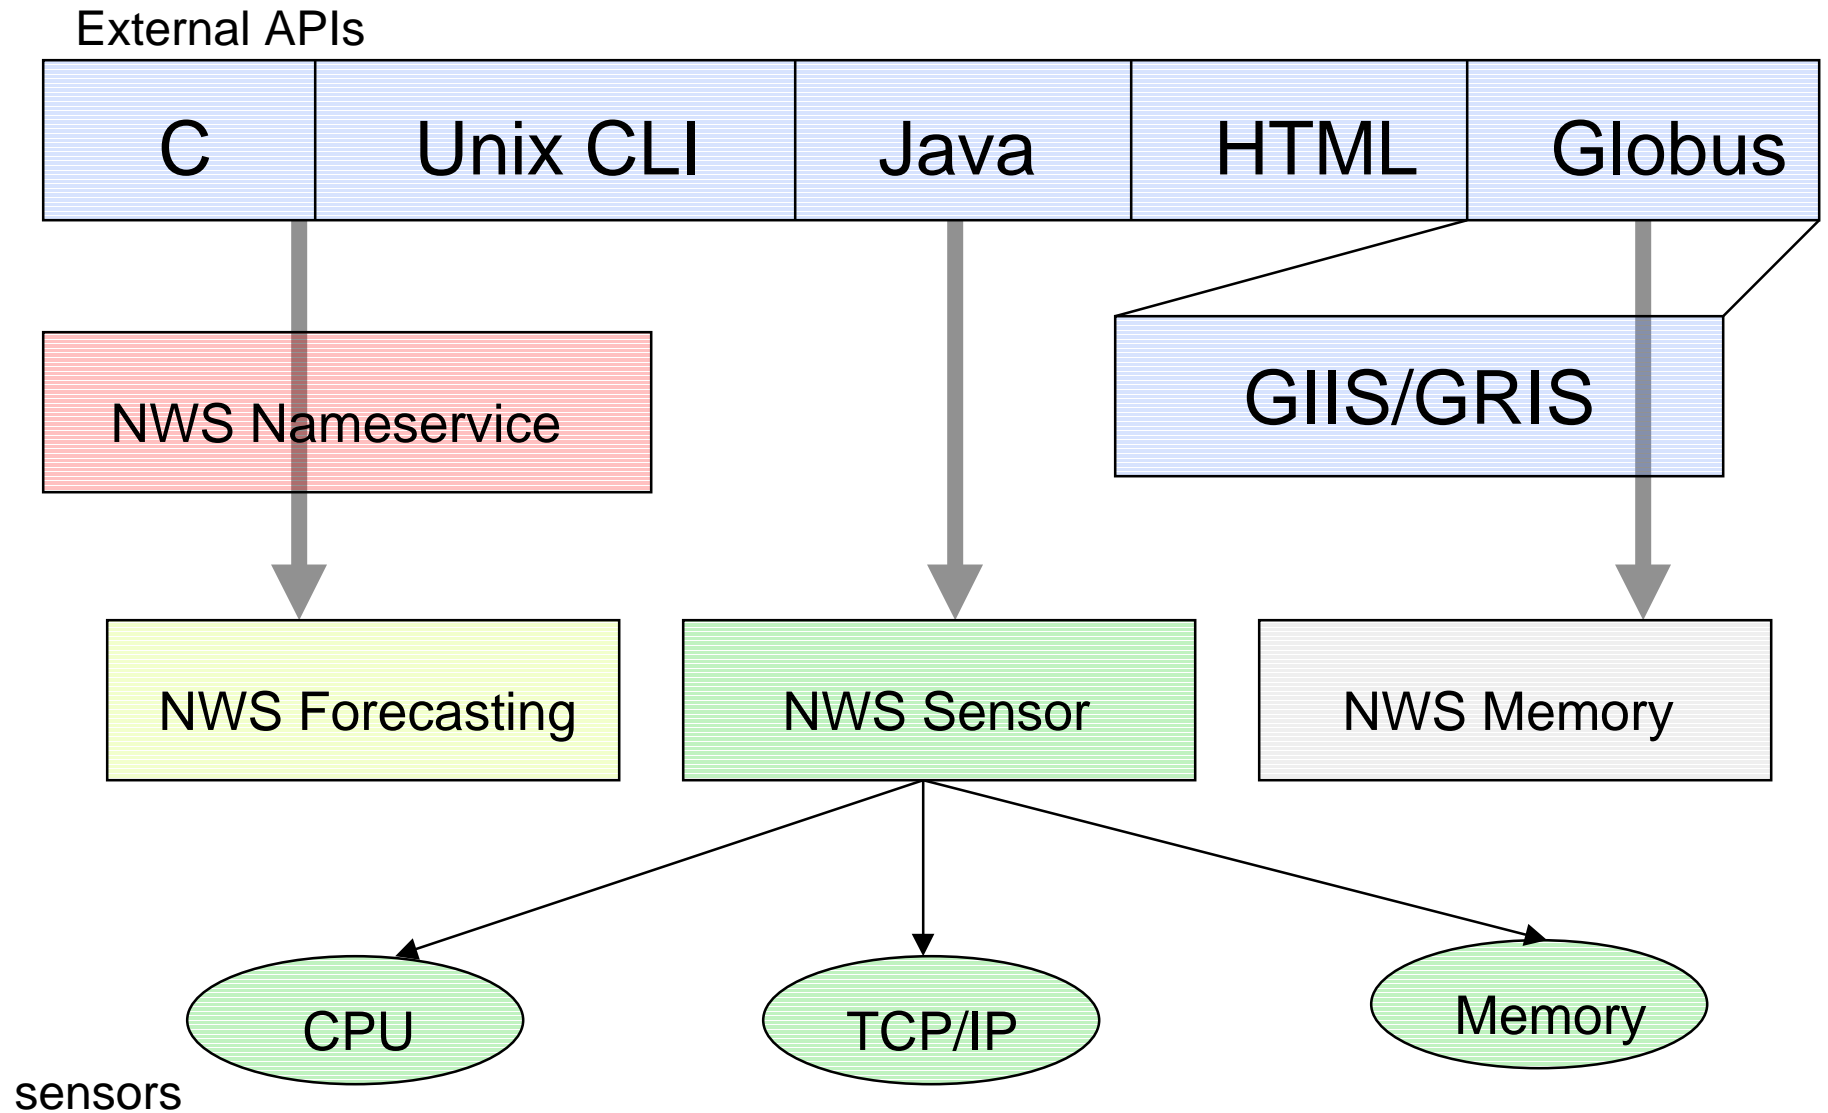
Archiving Data

- **NWS** persistent state subsystem is optimized to store time series
 - Circular buffer of recent data => **time-limited storage**
 - Default fetch order is **most-recent-first**
 - Performance optimized for **frequent single update** and **bulk-serial retrieval** (i.e. anti-LDAP)
- **Netlogger** is used to capture indefinite performance archives
 - ULM format (emerging standard)
 - Nifty visualization tools
 - XML?

What's Next?


- Active naming and the **GIIS/GRIS** architecture
 - **Goal**: to use the **Globus** GIIS/GRIS naming system in place of the NWS Nameservice
 - All **NWS** components will respond dynamically to LDAP queries
 - **Globus** users will be able to mix-and-match NWS components within the framework of the Grid Information Service
- Requires a **new DIT** structure for the Grid Information Service
 - Talks within Grid Forum cabal are underway

GIIS/GRIS In Place of NWS Nameservice




Fault Recognition

Network Fault Diagnosis Using Measurement Data




NWS Forecasting

Measurements and Forecasts



Using NWS Forecasts

Network Fault Diagnosis Using Forecast Data



The NWS Alarm System

- Users specify a resource and a performance range
- Trigger an alarm when forecast performance falls outside the specified range
- Both Java and Unix Daemons
 - Bells, whistles, flashing lights
 - email
- Useful as a *watchdog* for Grid operations

Status

- Globus interface to [NWS](#)
 - Version 1 up and functional
 - Version 2 almost ready for Beta testing
- [NWS](#) dynamic benchmarking library
 - Done
 - Used by AppLeS PST
- Netlogger archival support
 - Working but not complete
 - Beta test 10/15
- Globification of [NWS](#)
 - In design
- [NWS](#) Alarm system
 - Prototypes functioning
 - Needs work

People, Places, and Things

- Martin Swany -- swany@cs.utk.edu
 - Keeper of the software
 - NWS users mailing list
- <http://nws.cs.utk.edu>
 - Text being updated now
 - Interactive query page is functioning
- [Public Release 2.0 due out any day now](#)
 - Beta 7 released and in the hands of “friendly” users
- Thanks
 - Steve Fitzgerald -- ISI/USC
 - Graziano Obertelli -- UCSD
 - Patrick Geoffray – ENS Lyon
 - Chandra Krintz -- UCSD
 - Jim Hayes -- UCSD